

What Is Claimed Is:

1. A method of describing multiple packets to a communication apparatus with a single descriptor, the method comprising:
 - 5 receiving a header buffer comprising a header for each of N packets, wherein N>1;
 - receiving a data buffer comprising a payload for each of said N packets; in a single descriptor configured to be read by the communication apparatus, storing:
 - 10 a base address of said header buffer;
 - a base address of said data buffer;
 - a checksum start offset indicating where to compute a checksum on each of said N packets;
 - a checksum stuff offset indicating where to store the checksum in each of said N packets; and
 - 15 for each of said N packets:
 - a length of the payload for said packet;
 - a sample of the payload for said packet; and
 - a length of the header for said packet.
- 20
2. The method of claim 1, further comprising:
 - storing in said single descriptor:
 - a first indicator configured to indicate whether the headers for said N packets are packed in said header buffer; and
 - 25 a second indicator configured to indicate whether the payloads for said N packets are packed in said data buffer.

3. The method of claim 1, further comprising:
storing in said single descriptor, for each of said N packets, an offset of the
packet's header within said header buffer.
- 5 4. The method of claim 1, further comprising:
storing in said single descriptor, for each of said N packets, an offset of the
packet's payload within said data buffer.
- 10 5. The method of claim 1, wherein the communication apparatus is
configured for InfiniBand, the method further comprising:
storing in said single descriptor:
an R_key for said header buffer; and
an R_key for said data buffer.
- 15 6. A computer readable medium storing instructions that, when
executed by a computer, cause the computer to perform a method of describing
multiple packets to a communication apparatus with a single descriptor, the
method comprising:
receiving a header buffer comprising a header for each of N packets,
20 wherein N>1;
receiving a data buffer comprising a payload for each of said N packets;
in a single descriptor configured to be read by the communication
apparatus, storing:
a base address of said header buffer;
25 a base address of said data buffer;
a checksum start offset indicating where to compute a checksum on
each of said N packets;

a checksum stuff offset indicating where to store the checksum in each of said N packets; and

for each of said N packets:

a length of the payload for said packet;

5 a sample of the payload for said packet; and

a length of the header for said packet.

7. A processor-implementable method of using a single descriptor to facilitate the passing of multiple packets to a communication apparatus from a
10 device driver, the method comprising:

storing multiple packets' headers contiguously within a header buffer;

storing payloads of the multiple packets contiguously within a data buffer;

providing said header buffer and said data buffer to a device driver for a communication apparatus;

15 configuring a single descriptor to describe the multiple packets;

passing said single descriptor to the communication apparatus; and

at the communication apparatus, using said single descriptor to transmit the multiple packets.

20 8. The method of claim 7, further comprising:

configuring a second descriptor to reference said single descriptor, wherein said second descriptor is a traditional descriptor configured to describe a single packet;

passing the second descriptor to the communication apparatus; and

25 at the communication apparatus, reading the second descriptor to access said single descriptor.

9. The method of claim 7, wherein configuring a single descriptor comprises including in the single descriptor:

- a base address of said header buffer;
- a base address of said data buffer; and
- 5 for each packet in the multiple packets:
 - a length of a header of the packet; and
 - a length of a payload of the packet.

10. The method of claim 9, wherein configuring a single descriptor
10 further comprises including in the single descriptor:

- a checksum start value configured to indicate where, in each of the multiple packets, a checksum computation is to be initiated; and
- a checksum stuff value configured to indicate where, in each of the multiple packets, a checksum computation is to be stored.

15 11. The method of claim 9, wherein configuring a single descriptor further comprises including in the single descriptor:

- a packed header indicator configured to indicate whether the multiple packets' headers are packed in said header buffer; and
- 20 a packed data indicator configured to indicate whether the multiple packets' payloads are packed in said data buffer.

12. The method of claim 9, wherein configuring a single descriptor further comprises including in the single descriptor:

- 25 for each packet in the multiple packets, a sample of the packet's payload.

13. A computer readable medium storing instructions that, when

executed by a computer, cause the computer to perform a method of using a single descriptor to facilitate the passing of multiple packets to a communication apparatus from a device driver, the method comprising:

- 5 storing multiple packets' headers contiguously within a header buffer;
- storing payloads of the multiple packets contiguously within a data buffer;
- providing said header buffer and said data buffer to a device driver for a communication apparatus;
- configuring a single descriptor to describe the multiple packets;
- passing said single descriptor to the communication apparatus; and
- 10 at the communication apparatus, using said single descriptor to transmit the multiple packets.

- 14. A computer readable medium containing a data structure configured to describe multiple packets to a communication apparatus for transmitting the multiple packets, the data structure comprising:
 - a base address of a header buffer storing headers for the multiple packets;
 - a base address of a data buffer storing payloads for the multiple packets;
 - for each packet in the multiple packets:
 - a length of a header of the packet; and
 - a length of a payload of the packet;
 - a first indicator configured to indicate whether the headers are stored contiguously in said header buffer; and
 - a second indicator configured to indicate whether the payloads are stored contiguously in said data buffer.
- 20
- 25

- 15. The computer readable medium of claim 14, wherein the data structure further comprises:

a checksum start value configured to indicate where, in each of the multiple packets, a checksum computation is to be initiated; and

a checksum stuff value configured to indicate where, in each of the multiple packets, a checksum computation is to be stored.

5

16. A computer readable medium containing a data structure configured for describing multiple packets to a communication apparatus for transmitting the multiple packets, the data structure comprising:

a base address of a buffer storing multiple packets;

10 for each packet in the multiple packets, a length of the packet; and

a first indicator configured to indicate whether the packets are stored contiguously in said buffer.

17. The computer readable medium of claim 16, wherein the data structure further comprises, for each packet, if the packets are not stored contiguously in said buffer, an offset of the packet in said buffer.

18. An apparatus for transmitting packets, comprising:
a reader module configured to read a single descriptor configured to
20 describe multiple packets;
a retrieval module configured to retrieve the multiple packets; and
a transmitter module configured to transmit the multiple packets.

19. The apparatus of claim 18, wherein said retrieval module
25 comprises:

a header retriever configured to retrieve headers for the multiple packets;
and

20

a payload retriever configured to retrieve payloads for the multiple packets.

20. The apparatus of claim 18, wherein said single descriptor
5 comprises:

a base address of a header buffer storing headers for the multiple packets;

a base address of a data buffer storing payloads for the multiple packets;

for each packet in the multiple packets:

a length of a header of the packet; and

10 a length of a payload of the packet;

a first indicator configured to indicate whether the headers are stored contiguously in said header buffer; and

a second indicator configured to indicate whether the payloads are stored contiguously in said data buffer.

15

21. The apparatus of claim 18, wherein said single descriptor comprises:

a base address of a buffer storing multiple packets;

for each packet in the multiple packets, a length of the packet; and

20 a first indicator configured to indicate whether the packets are stored contiguously in said buffer.

22. A method of describing multiple packets to a communication apparatus, the method comprising:

25 for each packet, configuring a first data structure to identify:

a header length; and

a payload length;

storing in the first data structure headers for the packets;
configuring a first descriptor to identify the first data structure;
configuring said first descriptor to identify a location of the headers within
the first data structure; and
5 forwarding said first descriptor to the communication apparatus to
facilitate transmission of the packets.

23. The method of claim 22, further comprising:
10 storing in the first data structure payloads for a subset of the packets; and
configuring said first descriptor to identify a location of the payloads
within the first data structure.

24. The method of claim 22, further comprising:
15 storing in a second data structure payloads for a subset of the packets;
configuring a second descriptor to identify a location of the payloads
within the second data structure; and
forwarding said second descriptor to the communication apparatus.

25. The method of claim 22, wherein said configuring a first data
20 structure further comprises:
 configuring the first data structure to identify:
 a checksum start value configured to indicate where, in each of the
multiple packets, a checksum computation is to be initiated; and
 a checksum stuff value configured to indicate where, in each of the
multiple packets, a checksum computation is to be stored.
25

26. The method of claim 22, wherein said configuring a first data

structure further comprises:

configuring the first data structure to identify a type of checksum.

27. A computer readable medium containing data structures for
5 facilitating transmission of multiple packets from a communication apparatus, the
data structures comprising:

a first metadata structure configured to include:

a metadata section configured to identify, for each of the multiple
packets:

10 a header length; and

a payload length; and

a header section configured to store headers for the multiple
packets; and

a first descriptor configured to identify:

15 a memory location of said first metadata structure; and

a location of said header section within said first metadata
structure.

28. The computer readable medium of claim 27, further comprising:

20 a second metadata structure configured to store payloads for the multiple
packets; and

a second descriptor configured to identify a memory location of said
second metadata structure.

25 29. The computer readable medium of claim 27, wherein:

said first metadata structure is further configured to include:

a payload section configured to store payloads for the multiple

packets; and

said first descriptor is further configured to identify:

a location of said payload section within said first metadata structure.

5

30. The computer readable medium of claim 27, wherein:

said first metadata structure is further configured to include:

a checksum start value configured to indicate where, in each of the multiple packets, a checksum computation is to be initiated; and

10 a checksum stuff value configured to indicate where, in each of the multiple packets, a checksum computation is to be stored.

31. The computer readable medium of claim 27, wherein:

said first metadata structure is further configured to identify:

15 a type of checksum for checksumming the multiple packets.